Weakly Supervised Top-down Salient Object Detection

Hisham Cholakkal, Jubin Johnson, and Deepu Rajan

Abstract—Top-down saliency models produce a probability map that peaks at target locations specified by a task/goal such as object detection. They are usually trained in a fully supervised setting involving pixel-level annotations of objects. We propose a weakly supervised top-down saliency framework using only binary labels that indicate the presence/absence of an object in an image. First, the probabilistic contribution of each image region to the confidence of a CNN-based image classifier is computed through a backtracking strategy to produce top-down saliency. From a set of saliency maps of an image produced by fast bottom-up saliency approaches, we select the best saliency map suitable for the top-down task. The selected bottom-up saliency map is combined with the top-down saliency map. Features having high combined saliency are used to train a linear SVM classifier to estimate feature saliency. This is integrated with combined saliency and further refined through a multi-scale superpixel-averaging of saliency map. We evaluate the performance of the proposed weakly supervised top-down saliency against fully supervised approaches and achieve state-of-the-art performance. Experiments are carried out on seven challenging datasets and quantitative results are compared with 36 closely related approaches across 4 different applications.

Index Terms—Top-down saliency, salient object detection, weakly supervised training, object segmentation, semantic segmentation, object localization, object detection, CNN image classifier.

arXiv:1611.05345v2 [cs.CV] 17 Nov 2016

1 Introduction

The human visual system has the ability to zero-in rapidly onto salient regions in an image. Recently, there has been much interest among computer vision researchers to model this process known as *visual saliency*, which is attributed to the phenomenon of visual attention. It is beneficial in applications such as object detection/segmentation [1], [2], image retargeting [3] etc., since identification of salient regions reduces the search space for such high-level tasks. Salient regions in an image are indicated by a probability map called the saliency map. Fig. 1 shows saliency maps in the form of heat maps, where red indicates higher saliency.

In many instances, the salient region corresponds to a specific object in an image, in which case *salient object detection* becomes a more apt term, wherein pixels belonging to a salient object are assigned high saliency values. Broadly, there are two approaches to salient object detection: bottom-up (BU) [4] and top-down (TD) [5]. The feature contrast at a location plays the central role in BU salient object detection, with no regard to the semantic contents of the scene, although high-level concepts like faces have been used in conjunction with visual cues like color and shape [6]. The assumption that the salient object 'pops out' does not hold when there is little or no contrast between the object and the background. Furthermore, the notion of a salient object is not well-defined in BU models as seen in Fig. 1 (b, c, d) where recent methods [7], [4] and [8] show the potted plant in the background as salient to a user searching for the cat.

TD salient object detection is task-oriented and utilizes prior knowledge about the object class. For example, in semantic segmentation [9], a pixel is assigned to a particular object class, and a saliency map that aids in this segmentation must invariably be generated by a top-down approach. Fig. 1(e, f, g, h) show the saliency maps produced by the proposed method for person, cat, sofa and potted plant categories, respectively. TD saliency is also viewed as a focus-of-attention mechanism by which BU salient points that are unlikely to be part of the object are pruned [10].

Most methods for TD saliency detection learn object classes in a fully supervised manner using pixel-level labeling of objects [5], [11], [12]. Weakly supervised learning (WSL) alleviates the need for user-intensive annotation by utilizing only class labels for images. Moosmann *et al.* [13] propose a weakly supervised TD saliency method for image classification that employs iterative refinement of object hypothesis on a training image. Our method does not require any iterations, yet achieves state-of-the-art results compared to even fully supervised approaches [5], [14]. We first train a convolutional neural network (CNN) image classifier using image-level representation of CNN features, that gives a confidence score on the presence of an object in an image. The probabilistic contribution of each discriminative feature to this confidence score is represented in a TD saliency map, which is combined with a BU saliency map that is selected from several candidate BU maps through a novel selection strategy. Next, the saliency of each feature is separately evaluated using a dedicated feature classifier, as a means to assign non-zero saliency values to features from non-discriminative object regions, based on their dissimilarity with the background features. Saliency inference at a pixel involves combining the image classifier-based saliency map and the feature classifier-based saliency map.

A preliminary version of this paper was presented at CVPR 2016 [15]. The current version is revised with the following modifications: (i) sparse codes of SIFT features in [15] are replaced with CNN features; (ii) for a given task, a saliency-weighted max-pooling strategy is proposed to select a BU saliency map among several candidates, which is combined with TD saliency map; (iii) since CNN features span larger spatial neighborhood compared to SIFT features, contextual saliency in [15] is replaced with

• *H. Cholakkal, J. Johnson, and D. Rajan are with the School of Computer Science and Engineering, Nanyang Technological University, Singapore, 639798. E-mail: {hisham002, jubin001, asdrajan}@ntu.edu.sg.*

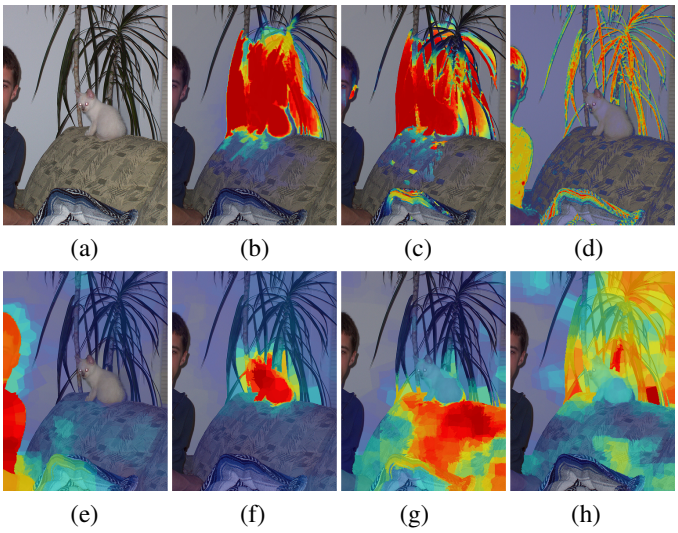

Fig. 1. Comparison of proposed top-down salient object detection with bottom-up methods. (a) Input image, bottom-up saliency maps of (b)MB [7], (c) MST [4], and (d) HC [8]; proposed top-down saliency maps for (e) person (f) cat (g) sofa and (h) potted plant categories.

CNN feature saliency; (iv) multi-scale averaging of saliency values within each superpixel is carried out to improve accuracy along object boundaries. These modifications lead not only to better performance than [15], but also with recent fully supervised TD approaches as shown in Fig. 2. Besides illustrating the accuracy of saliency maps produced by the proposed method, we also demonstrate the usefulness of our TD salient object detection for weakly supervised applications in semantic segmentation, object segmentation, object localization and object detection.

2 Related Work

We review related work in top-down saliency and relevant applications of CNN under weak supervision.

2.1 Top-down saliency frameworks

Kanan *et al.* [16] proposed a TD saliency approach which uses object appearance in conjunction with a location prior. Ineffectiveness of this prior largely affects the accuracy based on the position of the object within the image. Closer to our framework, Yang and Yang [5] proposed a fully supervised TD saliency model that jointly learns a conditional random field (CRF) and dictionary using sparse codes of SIFT features as latent variables. The inability to discriminate between objects having similar parts (e.g. wheels of car and motorbike) causes a large number of false detections. Kocak *et al.* [11] improved upon this by replacing SIFT features with the first and second order statistics of color, edge orientation and pixel location within a superpixel, along with objectness [17]. Although this improved the accuracy in distinguishing objects from background, it failed to discriminate between object categories, causing large number of false detections if the test image contained objects from other categories, as shown in Fig. 2 (b). Blocking artifacts are also observed in the saliency map at the superpixel boundaries because the superpixels are extracted on a single scale alone. Khan and Tappen [18] used label and location-dependent smoothness constraint in a sparse code formulation to produce a smooth saliency map compared to conventional sparse coding, but with additional computational cost. A joint framework for image classification and TD saliency is proposed in [19].

Zhu *et al.* [20] proposed a contextual-pooling based approach where LLC [21] codes of SIFT features are max-pooled in a local neighborhood followed by log-linear model learning. By replacing LLC codes with locality-constrained contextual sparse coding (LCCSC), Cholakkal *et al.* [12] improved on [20] with a carefully chosen category-specific dictionary learned from the annotated object area. Discriminative models [10], [22], [23] often represent a few patches on the object as salient and not the entire object. Hence, such models end up with low recall rates compared to [11], [24]. In [23], the task of image classification is improved using discriminative spatial saliency to weight visual features.

Recently, a fully supervised, CNN-based TD saliency method was proposed that utilized visual association of query images with multiple object exemplars [14]. They followed a two-stage deep model where the first stage learnt object-to-object association and the second stage learnt object-to-background discrimination. Each patch, extracted using a sliding window, is resized to 224×224 and input separately to the CNN. There are approximately 500 patches in an image of size 500×400, resulting in 500 forward passes through the network. Training the model required more than a week on a GPU. Our approach needs only one forward pass to extract CNN features for the entire image, which reduces the computation time significantly. It is still able to produce better saliency maps (Fig. 2 (f)) compared to [14] (Fig. 2 (e)). CNN-based saliency approaches [25], [26], [27], [28] learn category-independent salient features [29] from a large number of fully annotated training images [8]. Training or fine-tuning these saliency models [26], [28] took multiple days, even after initializing their models with convolutional filter weights pre-trained for image classification on ImageNet [30].

The use of weak supervision in TD saliency has largely been left unexamined. Gao *et al.* [10] used a weakly supervised setting where bottom-up features are combined with discriminative features that maximize the mutual information to the category label. Higher saliency values are assigned only to the features which are discriminative for a category in an image classification task, limiting its use in applications such as object segmentation, where all pixels of the object need to be identified accurately. In [13], a joint framework using classifier and TD saliency is used for object categorization by sampling representative windows containing the object. Their iterative strategy leads to inaccurate saliency estimation if the initialized windows do not contain the object.

2.2 CNN-based weakly supervised frameworks

Recently, CNN has been used in a number of weakly supervised object localization approaches [31], [32], [33], [34]. Multiple-instance learning is applied on CNN features in [33]. In [31], image regions are masked out to identify regions causing maximal activation. The outputs of CNN on multiple overlapping patches are utilized for object localization in [34]. All these approaches need multiple forward passes on a network to localize objects, which makes them computationally less efficient. Oquab *et al.* [32] applied global max-pooling to localize a point on objects. Global max-pooling is replaced by average pooling in [35] to help identify the full extent of the object as well. The underlying assumption is that the loss for average pooling enables the network to identify discriminative object regions. However, the spatial information is

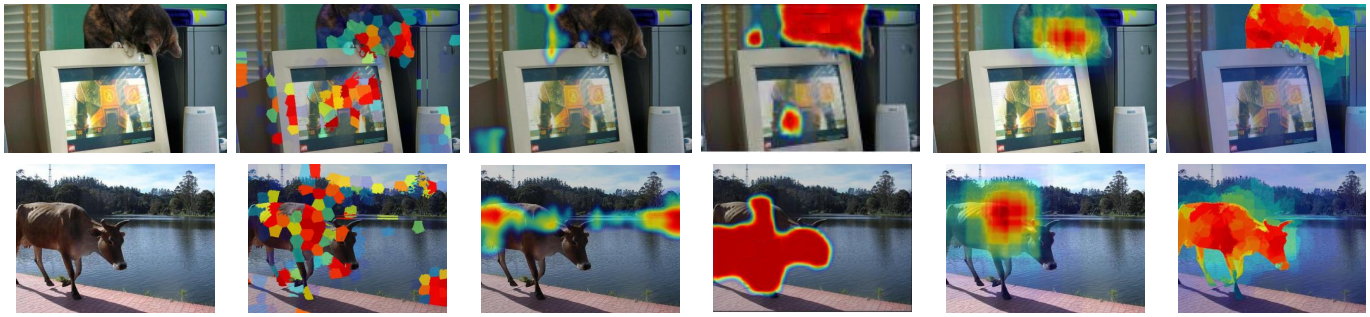

(a) (b) (c) (d) (e) (f)

Fig. 2. Visual comparison of the proposed weakly supervised approach with fully supervised top-down saliency approaches. (a) Input image, top-down saliency maps of (b) Kocak *et al.* [11], (c) LCCSC [12], (d) Yang and Yang [5], (e) Exemplar [14] and (f) proposed method for cat (top row) and cow (bottom row) categories.

lost, whereas it is retained in our framework via spatial pyramid pooling in the image classifier. The image classifier weights are reused for localization in [35]. We learn an additional feature classifier to better estimate saliency at non-discriminative object regions.

A weakly supervised, end-to-end CNN architecture is proposed in [36] for simultaneous object detection and image classification. Object detection requires classification of a large number of category-independent object proposals [37], [38]. On a test image, the CNN features are extracted on the original and flipped image at five scales totaling to 10 feature extraction iterations. In [39], category-specific object proposals are shown to be effective for weakly supervised object detection. Recent semantic segmentation, co-saliency and co-segmentation frameworks [9], [40], [41], [42] also train their CNN models in a weakly supervised setting.

Internal representations learned by CNN are visualized in [43], [44], [45], [46] for better understanding of its properties. [43] and [46] analyze the convolution layers using techniques such as deconvolutional networks. In [44] and [45], CNN features are inverted at different layers of the network including the fully connected layers, to analyze the visual encoding of CNN.

3 Proposed Method

In this section, we first describe the CNN image classifier and how the backtracking mechanism generates the TD saliency map. Next, a novel strategy to select a BU map is described followed by its integration with the TD saliency map to obtain combined saliency. A dedicated feature saliency model is learned on features extracted using this combined saliency to evaluate saliency at each feature independently. Finally, inference involves combining the two saliency maps to obtain the proposed TD map.

3.1 CNN Image Classifier

The CNN features of an image are extracted using VGG-16 [47] that has been pre-trained on the ImageNet ILSVRC 2012 data [30] with only image-level supervision (no bounding box annotation). All the layers of VGG-16 upto relu5_3 (third rectified linear unit in the fifth layer) are used for feature extraction and the convolution weights pre-trained for image classification are used without any fine-tuning. However, unlike [47], [48], we do not crop/resize the input image for feature extraction, but use the original image at its full resolution. A fixed-length, image-level representation of CNN features is obtained through a multi-scale spatial pyramid max-pooling as in Spatial Pyramid Pooling (SPP-net) [49] image classifier. We use a binary linear SVM after the spatial pyramid pooling layer, instead of fully connected layers in [49]. In order to reduce computations, we extract CNN features at a single image scale instead of multiple scales.

The spatial dimensions of an input image are down sampled by a factor of 16 at the relu5_3 feature map due to spatial max-pooling in the first four layers. There are 512 filter channels in conv5_3 (third convolution in the fifth layer), and hence each spatial location in relu5_3 feature map can be represented using a 512 dimensional feature vector, which we refer to as relu5 feature. The relu5 feature represents the overall response of multiple pixels from its receptive field in the original image.

Let $U = [u_1, u_2 ... u_m, ... u_M]$ denote M relu5 features each of dimension d. The spatial distribution of the features in the image is encoded in the spatial pyramid max-pooled image vector Z through a multi-scale max-pooling operation $F(u_1, u_2, ..., u_M)$ of the relu5 features on a 3-level spatial pyramid [50] as shown in Fig. 3. The i^{th} element z_i of Z is a max-pooled value derived using maximum operation on j^{th} elements of all relu5 features in a spatial pyramid region $\mathcal{R}$ defined by i, and $j = 1+(i-1)\ mod\ d$. i.e, $z_i = max\{u_{1j}, u_{2j}, ...u_{qj}\}, \forall\ 1, 2...q \in \mathcal{R}$.

Let the label $Y_k \in \{1, -1\}$ indicate the presence or absence of an object O in the k^{th} image. If $Y_k = 1$, it is a positive image, else it is a negative image. Image-label pairs (Z_k, Y_k) of T training images are used to train a binary linear SVM classifier [51], [52] with weight vector $W = [w_1, w_2 .... w_N]^\top$ and bias b. W is learnt separately for each object category. N is the length of the max-pooled image vector Z_k.

Given a validation/test image with max-pooled vector Z, the classifier score $W^\top Z + b$ indicates the confidence of the presence of object O in it, which is normalized to [0, 1] using the $sigmoid$ function. Our CNN image classifier is referred as cSPP in the following discussions.

3.2 Backtracking image classifier for saliency estimation

In the cSPP image classifier, both the linear-SVM and multi-scale max-pooling operations can be traced back to the relu5 feature locations. This enables us to analyze the contribution of each feature towards the final classifier score which is then utilized to generate the TD saliency map for an object.

First, we estimate the ability of a relu5 feature to represent a spatial pyramid region $\mathcal{R}$ by counting the number of times the

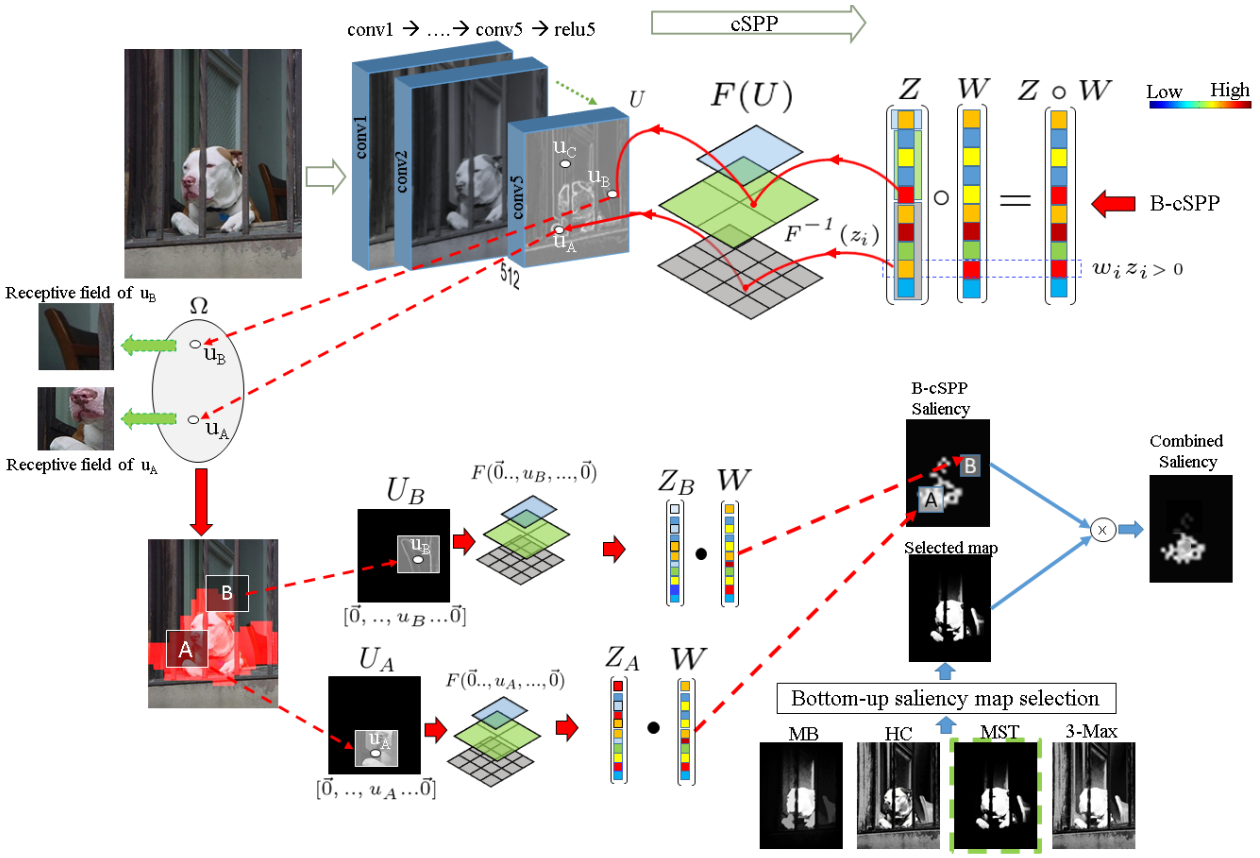

Fig. 3. Illustration of combined saliency estimation for *dog* category. Red arrows indicate the proposed backtracking strategy for top-down saliency (B-cSPP). From a set of BU saliency maps, the best one is selected and is integrated with B-cSPP saliency to produce combined saliency. 3-Max is the saliency map obtained by taking the maximum saliency at each pixel across the 3 BU saliency maps.

elements of that feature made it to the spatial pyramid (SP) max-pooled vector. We call this the representativeness, r_m, of a feature m. Representative features may either contribute positively or negatively to the classifier score with higher contribution indicating more relevance of the feature to an object O. The relevance of the feature to the object is denoted c_m.

For a positively contributing relu5 feature, it is possible that there are some negatively contributing elements among the total d elements. For example, let $[u_{m1}, 0, ...\ u_{mj}\ ...0, u_{md}]^\top$ be the m^{th} feature with its j^{th} element u_{mj} being a local maximum in its spatial pyramid region. Although u_{mj} contributes positively to the classifier confidence $W^\top Z + b$, the other non-zero elements u_{m1} or u_{md} may contribute negatively, indicating absence of the object. So, the relevance of a feature to the object requires its contribution to be computed in the absence of other features; this relevance is denoted p_m. The probability of a feature u_m belonging to an object, which in turn indicates the saliency of the object, depends on the three parameters - r_m, c_m and p_m as

$$P(r_m, c_m, p_m) = P(p_m|r_m, c_m)P(c_m|r_m)P(r_m). \quad (1)$$

The representative elements of the feature u_m are identified as

$$\Psi_m = \{i\delta(F^{-1}(z_i), u_{mj})\}, \quad \forall i \in \{1, 2, ..N\}, \quad (2)$$

where δ is the Kronecker delta function and F^{-1} is the inverse operation of spatial pyramid max-pooling and the location of z_i in Z identifies the region $\mathcal{R}$ in the spatial pyramid and its position j in the relu5 feature u_m. The probability of representativeness of the m^{th} feature to the image is then defined as

$$P(r_m) = \mathbf{card}(\Psi_m)/N, \quad (3)$$

where $\mathbf{card}(.)$ is the cardinality and N is the length of Z.

The linear SVM classifier confidence is a score indicating the presence of the object in the image, which increases from a definite absence ($score \leq -1$) to definite presence ($score \geq 1$). The confidence of $Y = 1$ is

$$\begin{aligned}
&\Theta\left(Y = 1 \,|\, F(u_1, u_2, ..., u_M)\right) \\
&\quad = W^\top F(u_1, u_2, ..., u_M) + b, \\
&\quad = W^\top Z + b \ = \sum_{\forall i \in \{1,..N\}} w_i z_i + b, \\
&\quad = \sum_{\forall i \in \Psi_m} w_i z_i \ + \sum_{\forall i \in \{1,..N\} \setminus \Psi_m} w_i z_i + b, \\
&\quad = \theta(c_m|r_m) \ + \sum_{\forall i \in \{1,..N\} \setminus \Psi_m} w_i z_i + b,
\end{aligned}$$

where $\theta(c_m|r_m)$ is the contribution of m^{th} feature to the image classifier confidence.

Given that the feature is representative of the image, the probability of it belonging to the object is

$$P(c_m|r_m) = \begin{cases} \beta\ (\theta(c_m|r_m)), & \text{if} \quad \theta(c_m|r_m) \geq 0, \\ 0, & \text{otherwise.} \end{cases} \quad (4)$$

where β is the sigmoid function that maps the confidence score to $[0, 1]$. Using the above probabilities, we select a set Ω of all features that contribute positively to the classifier confidence as

$$\Omega = \{P(c_m|r_m)P(r_m) > 0\}, \quad \forall m = 1, 2, ..., M. \tag{5}$$

The net contribution of a feature $m \in \Omega$ in the absence of other features is

$$P(p_m|r_m, c_m) = \beta\ (W^\top F(\vec{0}.., u_m, ..., \vec{0}) + b), \tag{6}$$

where $F(\vec{0}.., u_m, ..., \vec{0})$ is the spatial pyramid max-pooling operation performed by replacing all features except u_m with a zero vector $\vec{0}$ of size d to form max-pooled vector Z_m.

Implementation. Fig. 3 illustrates three relu5 features u_A, u_B and u_C. The confidence of the presence of object O in an image is indicated by the classifier score $W^\top Z + b$ as mentioned in the previous section and the element z_i of Z has a corresponding weight w_i. The elements from the Hadamard product $W \circ Z$ with $w_i z_i > 0$ mark the features u_A and u_B that contribute positively to the classifier confidence through a $F^{-1}(.)$ operation, i.e the set Ω. The contribution of feature u_A in the absence of other features is evaluated using max-pooling operations $F(\vec{0}.., u_A, .., \vec{0})$ in which all features except u_A are replaced with $\vec{0}$ forming max-pooled vector Z_A. The saliency of a feature m is given by

$$P(r_m, c_m, p_m) = \begin{cases} \beta\ (W^\top F(\vec{0}.., u_m, .., \vec{0}) + b) & \text{if } m \in \Omega, \\ 0 & \text{otherwise.} \end{cases} \tag{7}$$

Since this TD saliency of a feature is arrived at by backtracking the cSPP classifier, we call it *B-cSPP saliency* and the corresponding saliency map as *B-cSPP saliency map*. The feature u_A from the object (dog) region is assigned high B-cSPP saliency while u_B from background is assigned zero B-cSPP saliency.

3.3 Selection of bottom-up saliency map

As mentioned in Section 1, TD saliency helps prune out spurious BU saliency points to obtain a more accurate focus-of-attention [5], [10]. In this section, we propose a novel strategy to select the best saliency map from a set of BU maps based on saliency-weighted max pooling [19].

State-of-the-art BU saliency approaches [4], [7] can produce a category-independent saliency map for an image within 40 milliseconds. They assume image boundaries as the background while approaches such as [8] focus on feature contrast to estimate saliency. These approaches do not require any training and give reasonably good results. Since BU saliency maps are task-independent from a user's perspective, the definition of 'good saliency map' varies based on the application. For example, consider Fig. 1, where four different objects are present. If a user searches for a 'person' in the image, BU approaches [4], [7] that assume image boundary as the background fail to produce a 'good saliency map'. In such scenarios, an approach [8] that does not use such assumptions can produce better results. Thus, our objective is to develop a strategy to select a BU saliency method for a particular image that is best suited for the task at hand.

Our cSPP image classifier (W, b) which was trained to estimate the presence of object O in an image is employed to select a BU saliency map suitable for the task of identifying image regions that belong to object O. To achieve a one-to-one correspondence between pixels in the BU saliency map and the relu5 features, we downsample the saliency maps to the spatial resolution of feature map at relu5_3 , i.e, by a factor of 16. From n_ρ BU saliency maps, we need to select one for which features that belong to an object are assigned high saliency and those that do not belong to an object are assigned low saliency. For a max-pooled vector Z of an image, the SVM predicts a confidence score $W^\top Z + b$ which is proportional to the confidence of object presence in that image. i.e,

$$\begin{aligned} \Theta\ (Y = 1 \mid Z) &= W^\top Z + b = \sum_{\forall i \in \{1,..N\}} w_i z_i + b, \\ &= \sum_{\forall i \in\ \mathcal{I}^+} w_i z_i + \sum_{\forall i \in \mathcal{I}^-} w_i z_i + b. \\ \Theta\ (Y = 1 \mid Z) &= \sum_{\forall i \in \mathcal{I}^+} w_i z_i - \sum_{\forall i \in \mathcal{I}^-} |w_i| z_i + b \end{aligned} \tag{8}$$

where

$$\mathcal{I}^+ = \{i \mid w_i > 0\}, \quad \forall i \in \{1, 2, ..N\},$$
$$\mathcal{I}^- = \{i \mid w_i < 0\}, \quad \forall i \in \{1, 2, ..N\}.$$

Ideally, features belonging to object O contribute positively to the classifier confidence and hence they correspond to elements in Z whose indices belong to $\mathcal{I}^+$, while the background features result in $\mathcal{I}^-$ indices. It is to be noted that z_i is non-negative since it is derived from relu5 through max-pooling operation.

First, the m^{th} feature u_m is weighted with ρ_m^t, the BU saliency value for that feature estimated by t^{th} approach. i.e, $\hat{u}_m = u_m \times \rho_m^t$. The saliency-weighted relu5 features $\hat{U} = [\hat{u}_1, \hat{u}_2 ... \hat{u}_m, ... \hat{u}_M]$ are used to estimate the saliency-weighted max-pooled vector $\hat{Z}$ and similar to Eq. (8), the modified confidence score $\hat{B}(t) = \Theta\ (Y = 1 \mid \hat{Z})$ due to the t^{th} BU map is computed as,

$$\hat{B}(t) = \sum_{\forall i \in \mathcal{I}^+} w_i \hat{z}_i - \sum_{\forall i \in \mathcal{I}^-} |w_i| \hat{z}_i + b. \tag{9}$$

If higher values in the saliency map produced by algorithm t falls exactly on the object regions, the second summation will be largely reduced, due to weighting background indices with low saliency values and hence $\hat{B}(t)$ will be high. If some of the background also garners high saliency, then $\hat{B}$ will be relatively low. In order to reinforce the above assertion, we invert the saliency map (by subtracting saliency values from the maximum saliency value in the image), and recompute the saliency-weighted max-pooled vector $\tilde{Z}$ and $\tilde{B}(t)$ using same procedure. i.e,

$$\tilde{B}(t) = \sum_{\forall i \in \mathcal{I}^+} w_i \tilde{z}_i - \sum_{\forall i \in \mathcal{I}^-} |w_i| \tilde{z}_i + b. \tag{10}$$

If all object regions are assigned with higher saliency values in Eq. (10), higher weights are assigned to the background regions and lower weights to the salient regions, leading to a lower score of $\tilde{B}(t)$. Combining the above two observations, an ideal saliency map should maximize

$$\hat{B}(t) - \tilde{B}(t) = \sum_{\forall i \in \mathcal{I}^+} w_i (\hat{z}_i - \tilde{z}_i) - \sum_{\forall i \in \mathcal{I}^-} |w_i|\ (\hat{z}_i - \tilde{z}_i). \tag{11}$$

In order to prevent the selection of a map that assigns high saliency to the entire image, we impose a penalty of $1 - \mu_t$ on saliency map t with a mean saliency μ_t. Combining the above observations, the final objective function to select a BU saliency map is

$$\mathcal{B}(t) = \{ \sum_{\forall i \in \mathcal{I}^+} w_i(\hat{z}_i - \tilde{z}_i) - \sum_{\forall i \in \mathcal{I}^-} |w_i| \ (\hat{z}_i - \tilde{z}_i)\} \times (1 - \mu_t). \quad (12)$$

If the saliency map of t^{th} algorithm is not aligned with the object, then the false positives will increase $\hat{z}_i$ and decrease $\tilde{z}$ in $\mathcal{I}^-$, thus increasing the second term of Eq. (12). False negatives will reduce $\hat{z}_i$ and increase $\tilde{z}$ reducing the first term. Hence an inaccurate BU saliency map will result in low $\mathcal{B}(t)$. The saliency map that maximizes Eq. (12) is selected.

In addition to choosing individual BU saliency maps, we also analyze whether a combination of these maps has an effect on improving TD saliency. To this end, we combine saliency maps by picking the maximum saliency for each pixel and use Eq. (12) to select the best map from a set of saliency maps that includes the maximum map. In this section, we have assumed that the SVM weights learnt for an object is accurate and that the object appears only at locations where w_i are positive. Although this may not be always true, we retain this assumption since object locations are not available in a weakly supervised setting.

The B-cSPP saliency map and the selected bottom up saliency map are combined through a simple multiplication as shown in Fig. 3. We denote this combined saliency map as $\mathcal{H}$. Following [5], [10], [11], we also characterize our category-specific saliency inference framework as TD saliency even though there is a bottom-up component.

3.4 Feature saliency training

Image classifiers trained on image-level representation of features have shown to be effective in discriminative TD saliency estimation [10], [15], [23]. The combined saliency map $\mathcal{H}$ takes non-zero values only at discriminative image regions whose features make positive contribution to the image classifier confidence. The assumption is that the object appears only at grids in the spatial pyramid where w_i are positive, which may not be true across all images. Our objective is not limited to identifying the discriminative image regions, but to assign higher saliency values to all pixels belonging to the salient object. In order to independently estimate the saliency value of each relu5 feature, we also learn a top-down *feature saliency* model that uses a linear SVM learnt on positive and negative relu5 features from the training images. Since feature-level annotation is not available, we use object features extracted using the combined saliency map $\mathcal{H}$ to train the model.

From positive training images of object O, relu5 features with $\mathcal{H}$ saliency greater than 0.5 are selected as positive features with label $l = +1$. In order to prevent training features from non-discriminative object regions of positive images with negative label, only those features at which both B-cSSP and BU saliency are selected as negative features with label $l = -1$. Additionally, random features are selected from negative images with label $l = -1$. A linear SVM model with weight v and bias b_v is learned. Since the relu5 features are already computed for B-cSPP, learning of linear SVM is the only additional computation required to train this top-down model. The saliency map obtained from feature saliency is denoted $\mathcal{L}$.

3.5 Saliency inference

For inference on a test image, the combined saliency $\mathcal{H}$ and feature saliency are first integrated followed by multi-scale superpixel averaging and finally associated with the confidence of the image classifier to obtain the saliency at a pixel. While the combined saliency is obtained as described in Section 3.3, the feature saliency for a feature u_m is the probability of the feature belonging to an object computed by applying a sigmod function β to the linear SVM score,

$$P(l = 1 \mid u_m, v) = \beta \ (v^T u_m + b_v). \quad (13)$$

The feature saliency and combined saliency values are integrated using a mean operation to form the saliency map, $S_p = \frac{\mathcal{H}+\mathcal{L}}{2}$.

3.5.1 Multi-scale superpixel-averaging of saliency map

The low resolution saliency map S_p is upsampled to the original image size using bicubic interpolation. As a consequence, saliency values may not be uniform within a superpixel. Also, the saliency map will not be edge-aware with object regions spreading to the background. Hence, a multi-scale superpixel-averaging strategy is employed. The mean saliency at a superpixel (obtained by SLIC segmentation [53]) is assigned to every pixel in it. This process is repeated at multiple scales by varying the SLIC parameters. The resulting maps are averaged to produce a smooth, pixel-level saliency map S_{pix} that uniformly highlights the salient object and also produces a sharp transition at object boundaries.

3.5.2 Integrating with image classifier confidence

For a given image, the TD saliency map S_{pix} indicates the probable pixels that belong to object O. Since the presence of a specific object in a test image is not known apriori for applications such as semantic segmentation and object detection, the saliency map needs to be estimated for both positive and negative images. Hence, it is beneficial to integrate S_{pix} with a confidence score that indicates the presence of object O in at least one pixel in the image. For this, we use the same cSPP image classifiers learnt earlier for each category. The SVM associated with the cSPP image classifier gives a confidence score $\Phi(O)$ for a particular object O as $\Theta \ (Y = 1 \mid Z)$. These scores are scaled between 0 and 1 as

$$\hat{\Phi}(O) = \frac{exp(\Phi(O))}{\max\limits_{1 \le j \le n_c} \{exp(\Phi(j))\}}, \quad (14)$$

where, n_c is the total number of categories. Unlike soft-max that sums to 1, we normalize the score with the maximum because multiple categories can simultaneously appear in an image such as in PASCAL VOC-2012 [54]. In such scenarios, softmax will end up assigning a lower value to all positive categories. However, our objective is to identify the relative confidence across categories, and assign 1 to the most probable category. To reduce false detections from less probable categories, we assume values of $\hat{\Phi}(O)$ that are less than 0.5 as less important, and replace it with 0. This limits the number of probable object categories per image to less than 5 categories in most images, and hence the category-specific saliency map S_{pix} needs to be computed only for these few probable object categories. We compute the classifier-weighted, category-specific score for each object O,

$$S_{categ}(O) = S_{pix}(O) \cdot \hat{\Phi}(O). \quad (15)$$

3.5.3 Category-independent salient object detection

The proposed category-specific TD saliency map S_{categ} in Eq. (15) can be used to compute the category-independent saliency

value S_{ind}, by computing the maximum saliency value at each pixel (x,y) as

$$S_{ind}(x,\, y) = \max_{1 \le j \le n_c} \{S_{categ}(j)(x,\, y)\}.$$

Since the bottom-up information is integrated to S_{categ} through the combined saliency map $\mathcal{H}$, the $S_{ind}(x, y)$ gives an accurate estimate of saliency maps under free-viewing condition.

4 Applications

TD saliency [5], [11], [14], [15] and tasks like object detection, localization and segmentation mainly differ in their granularity of representation. Object detection produces a tight rectangular bounding box around all instances of objects belonging to user-defined categories. It is necessary to identify both the location as well as the extent of each object. The process of identifying the location of a particular object in an image, without marking the extent of the object, is referred to as object localization [32]. Object segmentation, also referred to as semantic object selection produces a binary mask with '1' indicating all pixels that belong to a user-defined object category. It differs from the task of semantic segmentation, where the objective is to classify each pixel in the image to one of predefined classes. In this section, we detail the use of our TD saliency framework for the above applications in a weakly supervised setting.

Semantic segmentation. The category-specific saliency maps in the proposed framework can be easily adapted for semantic segmentation. In the saliency map, a pixel with $S_{categ}(O) < 0.5$ is less likely to belong to an object O. The pixels at which the maximum saliency across all categories is less than 0.5 is more likely to be background. Hence, the additional map corresponding to the background category is generated as a uniform map with $S_{categ} = 0.5$. We assign to each pixel the category for which its saliency is the maximum.

Object segmentation. Conventional object segmentation approaches use scribbles or rectangular boxes to indicate the object of interest, while in our approach, only the semantic label of the object of interest is input to the system, similar to the semantic object selection [55]. We threshold our TD saliency map to identify definite foreground and background regions in an image, followed by Grab-cut [56] to accurately segment out the object of interest. Being a weakly supervised approach, framework is comparable to co-segmentation approaches that segment out a common object from a given set of images. We learn a model for the common object, which helps to achieve faster inference for a newly added test image, whereas co-segmentation approaches need to re-segment every image in the set upon encountering a new image.

Object localization. Object localization deals with locating object O within a positive image. Here, only the location of the object needs to be identified, not its extent. The peaks of our saliency map, S_{pix} indicates the location of object O,

$$Loc\,(O) = \underset{(x,\, y)}{\operatorname{argmax}}\, \{S_{Pix}(O)(x,\, y)\}.$$

Object detection. In object detection, multiple instances of the same object category need to be identified separately. This is more challenging than localization and especially so in a weakly supervised setting. Conventional object detectors such as R-CNN [1], [36] need to classify thousands of category-independent object proposals generated using selective search [37], [38]. This incurs a huge computational cost. The proposed TD saliency framework simplifies object detection by generating less than 5 proposals for an object category per image. First, the category-specific saliency $S_{categ}(O)$ is binarized by applying a threshold at 0.5. The smallest rectangular box enclosing each disconnected region is the detection box for object O. With this simple strategy, we achieve a performance which is comparable to dedicated weakly supervised object detectors [39].

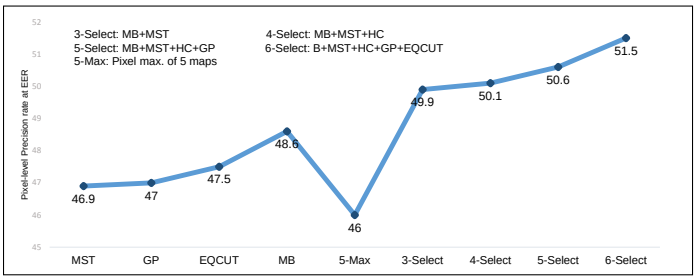

Fig. 4. Evaluation of selection of bottom-up approaches using pixel-level precision rate at EER (%) across 3 categories of Graz-02 dataset. The proposed selection strategy achieves better performance than the individual algorithms .

5 Experimental Evaluation

We evaluate our weakly supervised saliency model on PASCAL VOC-2012 [54], PASCAL VOC-2007 [57], Graz-02 [58] and PASCAL-S [29] datasets. Additionally, we use PASCAL VOC-2012 segmentation test set and validation set to evaluate weakly supervised semantic segmentation, Object Discovery dataset [59] to compare with semantic object selection and co-segmentation approaches, and validation set of PASCAL VOC-2012 detection challenge to evaluate object localization and object detection performance. In all these applications, we achieve state-of-the-art results compared to dedicated weakly supervised approaches.

Graz-02 dataset contains 3 object categories and a background category with 300 images per category. We split the images into training and testing sets following [24], i.e., 150 odd numbered images from each category are used for training and 150 even numbered images from each category for testing. PASCAL VOC-2012 is another challenging dataset with category-specific annotations for 20 object categories. It has 5717 training images and 5823 validation images for image classification/object detection challenge. There are 1464 training images, 1449 validation images and 1456 test images for segmentation challenge. Our PASCAL VOC-2012 saliency models are trained using 5717 training images for image classification task. There are 210 test images in the segmentation challenge of PASCAL VOC-2007.

PASCAL-S is a widely used dataset to evaluate category-independent saliency models. It has 850 images picked from the validation set of PASCAL VOC-2010 [60] segmentation images. Given the segmented objects in an image, the ground truth salient objects are marked by twelve subjects under free-viewing condition. We use Object Discovery dataset [59] to evaluate object segmentation. The dataset has three object categories, namely airplane, car and horse. Apart from 100 test images per category, there are 461, 1206 and 779 additional images for airplane, car and horse, respectively.

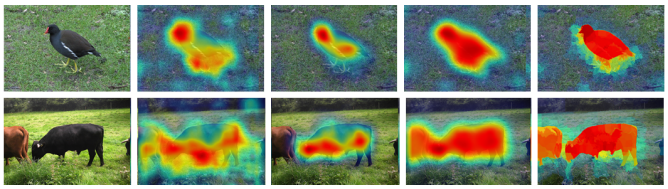

(a) (b) (c) (d) (e)

Fig. 5. Qualitative results at individual stages of the proposed method. (a) Input image, (b) B-cSPP saliency map, (c) (b) + bottom-up saliency, (d) (c) + feature saliency, (e) (d) + superpixel averaging.

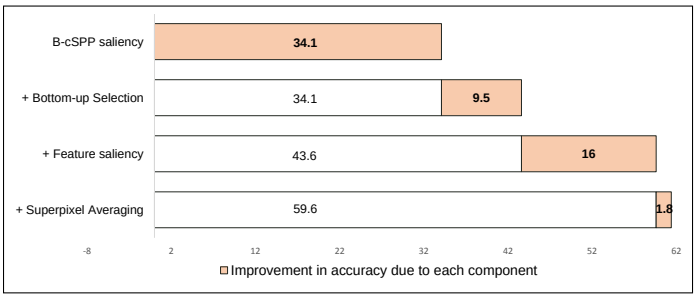

Fig. 6. Evaluation of individual stages of the proposed framework across 20 categories of PASCAL VOC-2012 using pixel-level precision rate at EER. The improvement in accuracy by the addition of each module is shaded.

5.1 Evaluation of selection of bottom-up saliency map

Fig. 4 illustrates the performance of the proposed strategy for selection of BU saliency map on positive test images, evaluated using mean of pixel-level precision rates at EER across all 3 object categories of Graz-02 dataset. Comparing the individual performances of 5 recent unsupervised algorithms HC [8], GP [61], EQCUT [62], MST [4] and MB [7] showed that MB outperforms others while HC has the lowest precision rate at EER. Since the Y-axis of the graph in Fig. 4 is limited to a range between 45 and 53, HC with mean precision rate at EER of 29.84% is not shown. Across individual categories, MB gives the best performance in bike category while EQCUT outperforms others in car and person categories.

First, we evaluate the performance of a maximum map formed by pixel-level maximum operation across the saliency maps of these 5 algorithms. Since the false positives from all the maps accumulate due to maximum operation, the mean precision rate at EER of this maximum map drops to 46% and denoted 5-Max in Fig. 4. Thus, combining BU maps without top-down information about the task can deteriorate the quality of the map.

Next, the proposed strategy to select the best saliency map among MB, MST and their maximum map is evaluated and shown as 3-Select in Fig. 4. Although MB outperforms MST in all the 3 categories, a performance boost of 1.3% is observed as a result of the selection of saliency map from MB, MST or maximum map for those images on which it outperforms others. The same procedure is repeated for the maps of MB, MST, HC and their maximum map and denoted 4-Select. In all the three categories, the performance of newly added HC algorithm is much lower than other approaches (less than 35%). We still observe an improvement of 0.2% in the mean precision rate at EER of 4-Select. Our approach automatically selects the best performing MB for 50.7% of the total images, 33.5% from MST, and only 8.5% from least performing HC. The remaining 7.3% are selected from the maximum map.

Similarly, addition of GP improved the accuracy by 0.5% in 5-Select. It is to be noted that MST, HC and GP are not the best performing algorithms in any of the individual categories, but their addition resulted in a gradual increase in the average accuracy. This shows that even though these algorithms have inferior performance in majority of the images in all 3 categories, they give better quality saliency maps for few images and the proposed selection strategy is able to accurately select those saliency maps.

Finally, 6-Select uses saliency maps of MB, MST, HC, GP, EQCUT, and the maximum map. In bike category, the largest number of maps are selected from MB (28% of bike images), which is the best performing algorithm for that category. Similarly, the largest number of car maps (23.3%) are selected from EQCUT, the best performing algorithm for car category. This shows the ability of the proposed strategy to carefully select the best algorithm for a given category.

Computation wise, GP and EQCUT take approximately 10 seconds per image to estimate saliency, while MB, MST and HC need less than 40 milliseconds. We use the latter three algorithms in our final framework (4-Select) to improve the runtime performance of the proposed method. The contribution of the selected BU saliency maps towards the final accuracy in PASCAL VOC-2012 dataset is analyzed in the next section.

5.2 Analysis of proposed framework

5.2.1 Contribution of individual modules

Fig. 5 shows the visual comparison of the effect of each stage in the proposed method. For the input images in Fig. 5(a), image regions containing bird's head and cow's legs make positive contribution to their image classifiers and are, therefore, assigned high B-cSPP saliency in Fig. 5(b). Combining B-cSPP saliency with BU saliency removed false detections in B-cSPP saliency as shown in Fig. 5(c). Integration of feature saliency assigns higher saliency value to the non-discriminative object regions (Fig. 5(d)). Finally, the addition of the multi-scale superpixel-averaging improved the accuracy along object boundaries as shown in Fig. 5(e).

We evaluate the improvement in the mean precision rate (%) at EER at each stage of our framework. The evaluation is done across 20 object categories of PASCAL VOC-2012 segmentation-validation set. The contribution of each component in the proposed saliency model to the final accuracy is shown shaded in Fig. 6. The accuracy of B-cSPP saliency is 34.1%. On adding the BU map to yield combined saliency, the accuracy increased to 43.6%, demonstrating the effectiveness of the proposed BU selection strategy.

Training the saliency model using negative patches from positive training images improve the accuracy by 5% in [15]. In the proposed framework, the accuracy of the combined saliency map $\mathcal{H}$ is improved by weighting B-cSPP with the selected BU map, which enabled us to train the feature saliency using negative patches from positive images. This resulted in an additional improvement of 3% in accuracy, totaling to 16% with the addition of feature saliency. This demonstrates that (i) learning a dedicated feature classifier plays an important role for TD saliency and (ii) combined saliency map $\mathcal{H}$ and feature saliency map complement each other. A similar trend in improvement was observed in [15], where a feature classifier is learnt using contextual max-pooled sparse codes. Since relu5 features span larger spatial

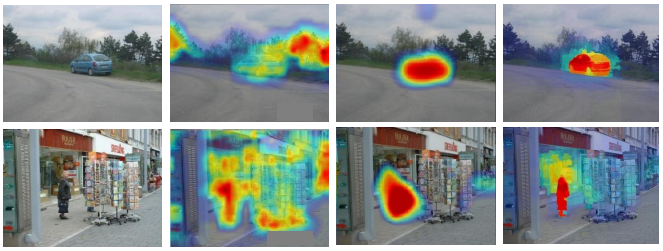

(a) (b) (c) (d)

Fig. 7. Visual comparison of the proposed method with other weakly supervised top-down saliency approaches. (a) Input image, top-down saliency maps of (b) DSD [10], (c) WS-SC [15] and (d) proposed method for car (top row) and person (bottom row) categories.

TABLE 1
Pixel-level precision rates at EER (%) on Graz-02.

Method	SV	Test set	Bike	Car	Person	Mean
1 - Yang and Yang [24]	FS	All test images	59.4	47.4	49.8	52.2
2 - Kocak *et al.* [11]	FS		59.9	45.2	51.5	52.2
3 - LCCSC [12]	FS		69.1	58.4	58.2	61.9
4 - CG-TD [19]	FS		64.4	50.9	56.4	57.2
5 - WS-SC [15]	WS		64.0	45.1	55.2	54.8
6 - Proposed	WS		80.5	61.4	75 .0	72.3
7 - MB [7]	US	Test images from respective category	54.67	39.03	52.04	48.58
8 - Aldavert *et al.* [65]	FS		71.9	64.9	58.6	65.13
9 - Fulkerson *et al.* [64]	FS		72.2	72.2	66.1	70.16
10 - Shape mask [63]	FS		61.8	53.8	44.1	53.23
11 - Yang and Yang [24]	FS		62.4	60	62	61.33
12 - Khan and Tappen [18]	FS		72.1	-	-	-
13 - CG-TD [19]	FS		67.3	59.8	57.1	61.4
14 - WS-SC [15]	WS		67.5	56.5	57.56	60.5
15 - Proposed	WS		84.1	81.5	81.8	82.5

context compared to SIFT features computed on 64×64 patches, contextual max-pooling on relu5 features is not required.

The feature saliency map $\mathcal{L}$ and combined saliency map $\mathcal{H}$ are integrated as $mean\ (\mathcal{H},\ \mathcal{L})$. Other combinations such as $max\ (\mathcal{H},\ \mathcal{L})$ and $\mathcal{HL} + 0.5(\mathcal{H} + \mathcal{L})$ gave similar results with less than 1% variation in accuracy. Taking the product of both saliency maps reduced the accuracy by 6% as the combined saliency $\mathcal{H}$ is often 0 in non-discriminative object regions and multiplication causes 0 values in such locations of the integrated map, disregarding feature saliency.

Finally, superpixel-averaging is applied at 6 different scales, by extracting 8, 16, 32, 64, 128 and 256 superpixels from an image. The saliency values at each pixel are further averaged across these 6 scales to get the saliency map S_{pix}. The quality of the saliency map at object boundaries is improved leading to 1.8% improvement in the accuracy, to obtain an accuracy of 61.4%. Since superpixel computation at multiple scales is time consuming relative to other modules, inference speed can be largely improved in applications such as object localization that do not require exact object boundaries by removing this step.

5.2.2 Performance comparison on different CNN architectures

We evaluated the performance of the proposed model on 5 different CNN architectures, namely VGG-F, VGG-M, VGG-S [48], VGG-16 and VGG-19 [47]. In all these architectures, layers from 'pool5' to 'prob' are removed to extract relu5 features. The performance is evaluated for all categories in PASCAL VOC 2012 dataset and mean of their pixel-level precision rate at EER is used to compare the architectures. The accuracy of saliency estimation varies by 10% across these architectures. The faster VGG-F gives the lowest accuracy (52%) and the deepest VGG-19 performs the best with 61.9%. The accuracy of VGG-M is 52.1%, followed by VGG-S (55%) and VGG-16 (61.4%). We do not extract relu5 features at multiple scales, nor do we crop or zero pad the input to a fixed size. Since there is no significant performance difference between VGG-16 and the deeper VGG-19, we use VGG-16 in our final framework across all datasets.

5.3 Comparison with other approaches

5.3.1 Graz-02 dataset

We report our pixel-level results on different test set configurations of Graz-02. First, the proposed saliency model is compared with other TD saliency algorithms [11], [12], [15], [19], [24] on all 600 test images. Second, for comparison with related approaches [63], [64], each object category is evaluated on test images from its respective category. Finally, to compare with [10], [16], results on 300 test images are evaluated, where 150 test images are from a single category and the remaining 150 are from the background.

The pixel-level comparisons in the first two test set configurations are shown in Table 1, where SV indicates supervision level with US, WS, FS referring to unsupervised, weakly supervised and fully supervised training, respectively. [11], [24] and [19] are fully supervised (FS), needing multiple iterations of CRF learning with sparse codes relearned at each iteration. Separate dictionaries are used for each object category. On the contrary, the proposed weakly supervised method does not require any iterative learning and the relu5 features are extracted with a single forward pass on the CNN. [15] does not require any iterative learning and uses a smaller dictionary of 1536 atoms, compared to 2048 atoms used in [12]. Despite incorporating objectness [17] and superpixel features to [24], the discriminative capability of [11] did not improve (row 2 vs row 1). The proposed weakly supervised method (row 6 and row 15) outperforms all other fully supervised TD saliency approaches [11], [19], [24] .

With respect to the second test configuration, [63] requires images to be marked as *difficult* or *truncated* in addition to the object annotation for training of shape mask. [65] uses 500,000 dictionary atoms in their fully supervised framework to obtain 65.13% (row 8), whereas the dimension of our relu5 feature is only 512. In this test setting, the proposed method achieves a mean accuracy of 82.5% outperforming the previous weakly supervised approach [15] by 22%.

As evident from Fig. 7(b), DSD [10] has limited capability to remove background clutter, resulting in poor performance of their model. DSD [10] and SUN [16] did not evaluate their model on Graz-02 dataset, but Yang and Yang [24] reported their patch-level precision rates at EER on 300 test images as 49.4% and 53.3%, respectively. Feature learning using independent component analysis helped SUN to perform better than DSD, but substantially poorer than [15] (65.4%) and the proposed method. It is to be noted that the performance of [10], [15], [16], [24] deteriorates while converting their patch-level results to pixel-level. The proposed weakly supervised method gives a mean pixel-level precision rate at EER of 73.1% which is better than the 70.16% and 70.49% reported by [11] and [12] respectively in this test setting. In all the three test settings, the proposed modifications enabled our current model to outperform [15] by more than 18% in accuracy, achieving state-of-the art performance. The use of CNN features contributes mainly to this performance boost. Qualitative comparisons with fully supervised and weakly supervised TD saliency approaches

TABLE 2
Pixel-level precision rates at EER on validation set of PASCAL VOC-2012 segmentation dataset.The proposed weakly supervised approach outperforms all fully supervised approaches including [14], which is based on CNN, in 14 out of 20 classes and in mean accuracy.

Method	SV	plane	bike	bird	boat	botl	bus	car	cat	chair	cow	table	dog	horse	moto	pers	plant	sheep	sofa	train	tv	Mean
Yang [24]	FS	14.7	28.1	9.8	6.1	2.2	24.1	30.2	17.3	6.2	7.6	10.3	11.5	12.5	24.1	36.7	2.2	20.4	12.3	26.1	10.2	15.6
Kocak [11]	FS	46.5	**45.0**	33.1	**60.2**	25.8	48.4	31.4	64.4	19.8	32.2	44.7	30.1	41.8	72.1	33.0	40.5	38.6	12.2	64.6	23.6	40.4
Exemplar [14]	FS	55.9	37.9	45.6	43.8	47.3	**83.6**	57.8	69.4	22.7	68.5	37.1	72.8	63.7	69.0	57.5	**43.9**	66.6	38.3	**75.1**	**56.7**	56.2
Oquab [32]	WS	48.9	42.9	37.9	47.1	31.4	68.4	39.9	66.2	**27.2**	54.0	38.3	48.5	56.5	70.1	43.2	42.6	52.2	34.8	68.1	43.4	48.1
Proposed	WS	**71.2**	22.3	**74.9**	39.9	**52.5**	82.7	**58.9**	**83.4**	27.1	**81.1**	**49.3**	**82.4**	**77.9**	**74.2**	**69.8**	31.9	**81.4**	**49.8**	63.2	53.3	**61.4**

TABLE 3
Precision rates at EER(%) on PASCAL VOC-2007.

Method	Yang and Yang [24]	LCCSC [12]	CG-TD [19]	WS-SC [15]	Proposed
Supervision	FS	FS	FS	WS	WS
Mean of 20 classes	16.7	23.4	23.81	18.6	42.1

are shown in Fig. 2 and Fig. 7, respectively.

5.3.2 PASCAL VOC-2012 segmentation dataset

In Table 2, we compare a recent CNN-based fully supervised TD saliency [14] with our method by evaluating on PASCAL VOC-2012 segmentation-validation set consisting of 1449 images. Similar to [14], each object category is evaluated only on positive images of that category. We did not fine-tune the convolution layers for this dataset, which took nearly 8 days on a GPU in [14]. Moreover, we only need a single CNN forward pass to extract features, while [14] requires 500 forward passes for an image. The presence of multiple, visually similar object classes in a single image is challenging for a weakly supervised approach. Inspite of this, we outperform the state-of-the art fully supervised approach [14] and the CNN-based weakly supervised object localization approach [32] in mean accuracy by 5% and 13%, respectively. We outperform [32] in 15 out of the 20 categories. The top-down selection of BU approach along with feature saliency plays an important role in this improved performance, especially in classes like aeroplane and sheep.

5.3.3 PASCAL VOC-2007 segmentation dataset

Following [11], [12] and [24], the saliency models are evaluated on 210 segmentation test images. We used the models trained on PASCAL VOC-2012 training set in this experiment. Separate sparse codes of size 512 are computed for each category in [11], [24] and [19]. [15] uses sparse coding on a common dictionary of 1536 atoms for all object classes. Similarly, a common feature code of $20 \times 512 = 10240$ elements is used in [12]. In our method, we compute 512 dimensional relu5 features which are common for all object categories.

Table 3 compares the pixel-level performance of the proposed WS method and patch-level results of FS top-down saliency approaches [12], [24] (these approaches did not report their pixel-level results on this dataset). We outperform [24], [19] and [12] in almost all categories and in mean precision rate at EER across 20 classes. A performance drop of 5 to 10% is reported by [66] while converting patch-level results of [24] to pixel-level, which further increases the performance gain of the proposed approach. Khan and Tappen [18] report pixel-level precision rates at EER only for cow category (8.5%) which is much lower than the proposed weakly supervised approach (52.3%).

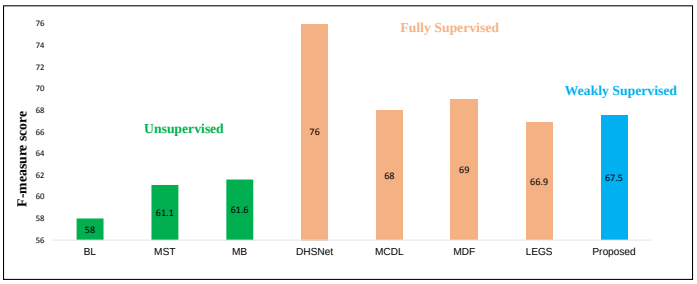

Fig. 8. Comparison of the proposed weakly supervised approach with state-of-the-art category-independent saliency approaches on PASCAL-S dataset. We achieve a performance comparable with deep learning-based fully supervised approaches.

5.3.4 Category-independent salient object detection

Category-independent saliency maps are obtained using the top-down models trained on PASCAL VOC-2012 training set through simple pixel-level maximum operation as explained in Section 3.5.3. Saliency values less than 0.5 are considered as background, and those between 0.5 and 1 are normalized to $[0, 1]$. The performance is evaluated on PASCAL-S dataset. Fig. 8 compares the proposed method against state-of-the-art category-independent approaches that include deep learning based fully supervised approaches such as MCDL [26], LEGS [27], MDF [28] and DHSNet [25]. The performance metric, F-measure is

$$f_\eta = \frac{(1+\eta^2) \cdot Precision \cdot Recall}{\eta^2 \cdot Precision + Recall},$$

where $\eta^2 = 0.3$ [29]. Following [25], precision and recall are computed by binarizing each saliency map at an image adaptive threshold, which is twice the average value of the saliency map.

The proposed weakly supervised method achieves an f-measure of 67.5, which is comparable with fully supervised LEGS, MCDL and MDF. We use only 5717 images from PASCAL VOC-2012 training set, which is much smaller compared to the training data used by fully supervised approaches shown in Fig. 8. For example, DHSNet uses nearly 10,000 fully annotated images from multiple datasets such as MSRA 10K [69] and DUT-OMRON [70]. Data augmentation is used to further increase the number of training images. With less supervision and lesser training data, we achieve a performance comparable with these fully supervised approaches. Qualitative results are present in the supplementary material.

5.4 Computation time

Training of the proposed framework is significantly faster than [11], [24] and [14] since we do not use iterative dictionary learning or fine-tuning of convolutional filter weights. MATLAB implementations of all approaches were evaluated on a PC running on Intel Xeon 2.4 GHz processor. Despite a non-parallelized

TABLE 4
Intersection over union (IOU) for semantic segmentation on validation set and test set of PASCAL VOC-2012.

Method	BG	plane	bike	bird	boat	botl	bus	car	cat	chair	cow	table	dog	horse	moto	pers	plant	sheep	sofa	train	tv	Mean
										Val												
MIL-FCN [67]	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	24.9
MIL-Base with ILP [9]	73.2	25.4	18.2	22.7	**21.5**	28.6	39.5	44.7	46.6	11.9	40.4	11.8	45.6	40.1	35.5	35.2	20.8	41.7	17.0	34.7	30.4	32.6
EM adapt [68]	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	38.2
Pathak [40]	68.5	25.5	18.0	25.4	20.2	36.3	46.8	**47.1**	48.0	**15.8**	37.9	**21.0**	44.5	34.5	46.2	40.7	**30.4**	36.3	**22.2**	38.8	**36.9**	35.3
Proposed	**77.7**	**57.4**	**18.6**	**58.2**	19.3	**40.5**	**62.0**	40.9	**69.7**	11.0	**50.7**	14.3	**65.9**	**49.3**	**50.9**	**54.8**	13.5	**54.2**	21.5	**47.0**	36.2	**43.5**
										Test												
Pathak [40]	68.5	25.5	18.0	25.4	20.2	36.3	46.8	**47.1**	48.0	**15.8**	37.9	21.0	44.5	34.5	46.2	40.7	**30.4**	36.3	22.2	38.8	**36.9**	35.3
EM adapt [68]	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	39.6
Proposed	**79.6**	**58.2**	**25.6**	**65.4**	19.2	**44.3**	**60.2**	39.5	**62.1**	10.6	**45.5**	**22.7**	**65.2**	**50.7**	**56.5**	**53.7**	14.3	**51.8**	**24.9**	**40.8**	34.9	**44.1**

Fig. 9. Semantic segmentation using our TD saliency map. Input image, semantic segmentation result produced by our framework and the ground truth for semantic segmentation are shown in adjacent columns.

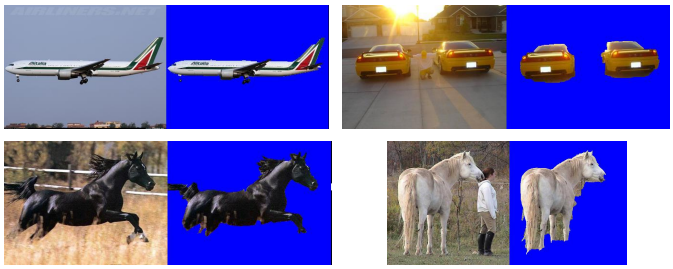

Fig. 10. Object segmentation using our top-down saliency map. Input image and the object segmentation results produced by our framework are shown in adjacent columns.

implementation, our approach requires 1 hour to train all 3 object categories of Graz-02. This is significantly faster compared to [24] which takes 4 hours and 49 minutes, and also to [11], which takes 30 hours and 10 minutes. The fully supervised algorithm proposed in [19] also needs 3 hours and 34 minutes to train Graz-02 models. Similarly, the proposed saliency models for all 20 categories of PASCAL VOC-2012 are trained within 6 hours and 20 minutes. Inspite of a parallel execution using GPU, [14] needs 8 days to train their model on PASCAL VOC-2012.

On a Graz-02 test image of size 640×480 pixels, our unoptimized, non-parallel MATLAB execution took 8 seconds for inference. The multi-scale superpixel averaging step takes up a major part of this running time, which can be removed at the expense of a 2% reduction in accuracy. With similar settings, [24] needs 5.5 seconds and [11] needs 28 seconds for inference. Since all the saliency models share common relu5 features, it reduces the computational time of the proposed method while inferring multiple saliency models on the same test image. However, [24] and [11] calculate sparse codes for each model separately and hence, the inference pipeline needs to be repeated for each category on a test image. This requires multi-fold inference time on images with multiple categories. On PASCAL VOC-2012 test images of size 500×350 pixels, our parallel MATLAB execution without GPU took an average inference time of 1.8 seconds as compared to 4 seconds in [14].

5.5 Applications

5.5.1 *Weakly supervised semantic segmentation*

In PASCAL VOC-2012 semantic segmentation task, each pixel in the image needs to be classified to one of 21 categories comprising background and 20 object categories. The proposed approach achieves state-of-the-art performance on both validation set and test set of PASCAL VOC-2012 segmentation challenge as shown in Table 4. The segmentation results on the validation set and test set are denoted Val and Test, respectively. Our results are reported from PASCAL VOC-2012 evaluation server which uses intersection over union (IoU)) as the evaluation metric. We outperform [40] in 14 out of 21 classes in the validation set and by nearly 8% in the mean IOU in both validation set and test set. A larger training set of 10582 and 12000 images are used by [40] and [68], respectively. It comprises additional images collected by Hariharan *et al.* [71] along with PASAL VOC 2012 training set. We trained our models using 5717 images from PASCAL VOC-2012 training set.

Fig. 9 compares the qualitative results obtained by the proposed method against the ground truth. Majority of the horse pixels are classified accurately in row 2, despite the presence of person. Similarly most of the person pixels are classified correctly, despite the size of person being small compared to motorbike in the top row. We do not use Grab-cut [56] or similar energy minimization techniques for semantic segmentation. The class label for each pixel is produced by simple maximum operation on the TD saliency maps.

5.5.2 *Weakly supervised object segmentation*

Object segmentation accuracy is evaluated on 100 test images from each category of Object Discovery dataset [59]. 300 images from each category are used to train our saliency model, along with 300 negative images from Graz-02 dataset. Qualitative results are shown in Fig. 10. Multiple instances of car are accurately segmented out as shown in the row 1. In row 2, the proposed approach could accurately segment out the horse. Quantitative comparisons with state-of-the-art co-segmentation approaches are shown in Table 5. The Jaccard similarity, i.e, intersection over union (IOU) with the ground-truth is evaluated as in [59]. In all the three categories, we achieve state-of-the-art performance compared to related co-segmentation [41], [59] and co-saliency [42] approaches. The semantic object selection [55] uses additional supervision by collecting positive training images with white

TABLE 6
Average precision of object localization on PASCAL VOC-2012 detection validation set.

Method	SV	plane	bike	bird	boat	botl	bus	car	cat	chair	cow	table	dog	horse	moto	pers	plant	sheep	sofa	train	tv	Mean
										Exact												
RCNN [1]	FS	86.5	72.1	74.2	66.7	43.1	78.3	68.8	80.8	44.9	62.3	51.1	74.4	73.6	83.0	83.0	49.2	78.4	40.6	74.1	69.2	67.7
Exemplar [14]	FS	86.8	87.2	72.7	46.8	31.7	91.0	58.6	95.2	44.5	94.8	41.5	87.0	91.4	**94.3**	**89.2**	57.7	**93.5**	59.2	84.7	60.5	73.4
ProNet [39]	WS	89.4	78.1	79.2	73.7	39.9	84.2	61.2	86.4	42.1	67.7	53.2	84	81.8	82.3	84.1	39.2	81.9	48.2	80.8	58.1	69.8
ProNet + classifier [39]	WS	90.3	82	81.2	**75.4**	44.5	84.9	62.5	89.8	47.2	78.9	55.6	88.1	84.3	85.7	85.6	44.4	84.3	50	81.8	65.1	73.1
Proposed	WS	**97.6**	**93.1**	**95.6**	73.4	**60.6**	**91.8**	**77.9**	**97.4**	**52.0**	**80.8**	**84. 1**	**96.6**	**91.4**	91.6	71.9	**62.4**	65.1	**84.4**	**92.5**	**84.1**	**82.2**
										18 Pix.												
RCNN [1]	FS	92.0	80.8	80.8	73.0	49.9	86.8	77.7	87.6	50.4	72.1	57.6	82.9	79.1	89.8	88.1	56.1	83.5	50.1	81.5	76.6	74.8
Fast RCNN [76]	FS	95.2	88.2	88.4	77.9	49.0	93.4	83.6	95.1	59.4	86.6	71.0	92.6	93.1	93.0	92.2	58.2	88.0	63.6	91.9	77.3	81.9
Oquab *et al.* [32]	WS	90.3	77.4	81.4	79.2	41.4	87.8	66.4	91.0	47.3	83.7	55.1	88.8	93.6	85.2	87.4	43.5	86.2	50.8	86.8	66.5	74.5
ProNet [39]	WS	91.6	82	85.1	78.6	45.9	87.9	67.1	92.2	51	72.9	60.8	89.3	85.1	85.3	86.4	45.6	83.5	55.1	85.6	65.9	74.8
ProNet + box classifier [39]	WS	92.6	85.6	87.4	**79.6**	48.3	88.7	68.9	94.2	54.6	83.2	62.8	92.0	89.9	88.2	87.1	49.2	86.9	57.2	86.8	70.0	77.7
Proposed	WS	**98.2**	**94.1**	**96.0**	79.1	**66.2**	91.8	79.5	**98.1**	58.5	85.4	**87.4**	**96.8**	92.0	**93.8**	76.0	**64.5**	69.3	**86.7**	**93.4**	**88.0**	**84.7**

TABLE 5
Comparison of proposed weakly supervised approach with object segmentation approaches on Object Discovery dataset, evaluated using Jaccard similarity.

Method	Airplane	Car	Horse	Mean
Joulin *et al.* [72]	15.4	37.2	30.2	27.6
Joulin *et al.* [73]	11.7	35.2	29.5	25.5
Kim *et al.* [74]	7.9	0.04	6.43	4.79
Object Discovey [59]	55.8	64.4	51.6	57.3
Koteshwar *et al.* [75]	56	69	55	60
Zhang *et al.* [42]	53.5	58.8	52.2	54.8
Quan *et al.* [41]	56.3	66.8	58.1	60.4
WS-SC [15]	57.3	67.4	50.51	58.4
Object selection [55]	64.3	71.8	55.1	63.7
Proposed	**65.0**	**77.3**	**61.6**	**68.0**

Fig. 11. Object localization using our top-down saliency map.

background using an internet search. Inspite of this modification, they could only achieve an average accuracy of 63.73%, which is lower than our mean accuracy of 68.0% across 3 categories.

5.5.3 *Weakly supervised object localization*

Presence of multiple objects in an image makes object localization on PASCAL VOC-2012 detection set a challenging task, especially in a weakly supervised setting. The location of the maximum value in the TD saliency map of an object category is used for its localization as explained in Section 4. Since an accurate estimate of object boundaries are not required, we replaced the multi-scale superpixel averaging with an averaging filter on a rectangular window of size 64×64 pixels for faster inference. The location that falls exactly within any ground truth bounding box associated for a given category is assumed correct and the average precision is calculated as in [14]. In [32], average precision is evaluated by giving an error tolerance of 18 pixels to the predicted location. We evaluated our model in both these settings denoted Exact and 18 Pix and corresponding results are compared with state-of-the art approaches as shown in Table 6. In both the evaluation settings, we achieve a performance which is comparable to fully supervised TD saliency approaches and dedicated object detectors such as faster RCNN [77]. Fig. 11 shows some qualitative results obtained using the proposed method in localizing multiple objects. Partially occluded objects such as motorbike and car are localized accurately despite the presence of other distracting objects.

TABLE 7
Comparison with weakly supervised object detection approaches on PASCAL VOC-2012 validation dataset, measured by average precision.

Method	Oquab *et al.* [32]	ProNet [39]	ProNet+Classifier [39]	Proposed
mAP (Mean of 20 Classes)	11.74	13	15.5	**20.4**

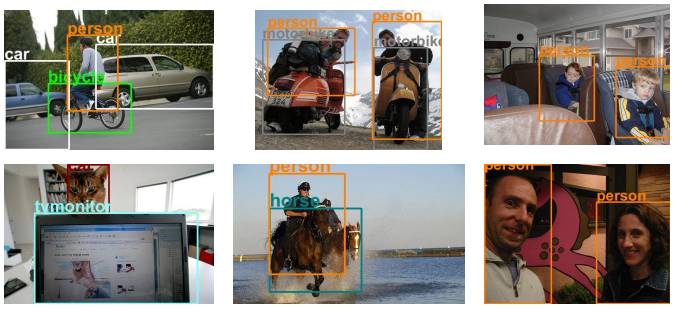

Fig. 12. Object detection using our top-down saliency map.

5.5.4 *Weakly supervised object detection*

The object detection boxes produced by a simple binarization of our saliency maps is shown to be comparable with dedicated weakly supervised object detectors in Table 7. We outperform [39] which uses an additional box classifier to classify their object proposal boxes. We consider all category-specific object boxes as positive detections. PASCAL VOC 2012 evaluation server is used to estimate object detection accuracy, where a detection having an $IOU > 0.5$ with the ground truth rectangular bounding box is considered as true positive. In Fig. 12, multiple overlapping objects are accurately detected by the proposed strategy. Multiple instances of person, motorbike and car are also detected. The horse and bicycle are accurately detected despite the presence of other categories in the image. Similarly, an accurate bounding box around the cat is marked in an image that also contains a TV monitor.

Limitations. Similar to other weakly supervised approaches [36], [39], the proposed approach has limited ability to discrim-

inate among multiple instances of an object which are spatially adjacent. This causes low performance for object detection, compared to state-of-the-art fully supervised object detectors [78]. Examples are provided in the supplementary material.

6 Conclusion

In this paper, a CNN feature-based weakly supervised salient object detection approach is proposed. A novel strategy to select a BU saliency map that suits a top-down task is proposed. Contribution of relu5 features at different spatial locations are estimated to compute a novel B-cSPP saliency. The top-down B-cSPP saliency is integrated with the BU saliency map and produces a combined saliency which is further integrated with feature saliency. The proposed weakly supervised top-down saliency model achieves state-of-the-art performance in top-down salient object detection by outperforming even fully supervised CNN-based approaches. Moreover, the top-down saliency maps of different object categories are combined to produce a category-independent saliency map that can estimate salient objects under free-viewing condition. Finally, through quantitative comparisons, we demonstrated the usefulness of proposed saliency map for four different applications. We plan to extend our framework to videos for weakly supervised salient object detection.

References

[1] R. Girshick, J. Donahue, T. Darrell, and J. Malik, "Rich feature hierarchies for accurate object detection and semantic segmentation," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2014, pp. 580–587.

[2] Y. Jia and M. Han, "Category-independent object-level saliency detection," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2013.

[3] D. Vaquero, M. Turk, K. Pulli, M. Tico, and N. Gelfand, "A survey of image retargeting techniques," in *SPIE Optical Engineering + Applications*, 2010, pp. 779–814.

[4] W.-C. Tu, S. He, Q. Yang, and S.-Y. Chien, "Real-time salient object detection with a minimum spanning tree," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, June 2016.

[5] J. Yang and M.-H. Yang, "Top-down visual saliency via joint crf and dictionary learning," *IEEE Trans. Pattern Anal. Mach. Intell.*, 2016.

[6] T. Liu, Z. Yuan, J. Sun, J. Wang, N. Zheng, X. Tang, and H.-Y. Shum, "Learning to detect a salient object," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 33, no. 2, pp. 353–367, 2011.

[7] J. Zhang, S. Sclaroff, Z. Lin, X. Shen, B. Price, and R. Měch, "Minimum barrier salient object detection at 80 fps," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2015.

[8] M. M. Cheng, N. J. Mitra, X. Huang, P. H. S. Torr, and S. M. Hu, "Global contrast based salient region detection," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 37, no. 3, pp. 569–582, March 2015.

[9] P. O. Pinheiro and R. Collobert, "Weakly supervised semantic segmentation with convolutional networks," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, vol. 2, no. 5, 2015, p. 6.

[10] D. Gao, S. Han, and N. Vasconcelos, "Discriminant saliency, the detection of suspicious coincidences, and applications to visual recognition," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 31, no. 6, pp. 989–1005, 2009.

[11] A. Kocak, K. Cizmeciler, A. Erdem, and E. Erdem, "Top down saliency estimation via superpixel-based discriminative dictionaries," in *Proc. British Mach. Vis. Conf.*, 2014.

[12] H. Cholakkal, D. Rajan, and J. Johnson, "Top-down saliency with locality-constrained contextual sparse coding," in *Proc. British Mach. Vis. Conf.*, 2015.

[13] F. Moosmann, E. Nowak, and F. Jurie, "Randomized clustering forests for image classification," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 30, no. 9, pp. 1632–1646, Sept 2008.

[14] S. He, R. W. Lau, and Q. Yang, "Exemplar-driven top-down saliency detection via deep association," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2016.

[15] H. Cholakkal, J. Johnson, and D. Rajan, "Backtracking scspm image classifier for weakly supervised top-down saliency," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2016.

[16] C. Kanan, M. H. Tong, L. Zhang, and G. W. Cottrell, "Sun: Top-down saliency using natural statistics," *Vis. cogn.*, vol. 17, no. 6-7, pp. 979–1003, 2009.

[17] B. Alexe, T. Deselaers, and V. Ferrari, "What is an object?" in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2010.

[18] N. Khan and M. F. Tappen, "Discriminative dictionary learning with spatial priors." in *Proc. Int. Conf. Image Proc.*, 2013.

[19] H. Cholakkal, J. Johnson, and D. Rajan, "A classifier-guided approach for top-down salient object detection," *Signal Process. Image Commun.*, vol. 45, pp. 24–40, 2016.

[20] J. Zhu, Y. Qiu, R. Zhang, J. Huang, and W. Zhang, "Top-down saliency detection via contextual pooling," *J. Signal Process. Syst.*, vol. 74, no. 1, pp. 33–46, 2014.

[21] J. Wang, J. Yang, K. Yu, F. Lv, T. Huang, and Y. Gong, "Locality-constrained linear coding for image classification," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2010.

[22] B. Yao, A. Khosla, and L. Fei-Fei, "Combining randomization and discrimination for fine-grained image categorization," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2011.

[23] G. Sharma, F. Jurie, and C. Schmid, "Discriminative spatial saliency for image classification," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2012, pp. 3506–3513.

[24] J. Yang and M.-H. Yang, "Top-down visual saliency via joint crf and dictionary learning," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2012.

[25] N. Liu and J. Han, "Dhsnet: Deep hierarchical saliency network for salient object detection," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, June 2016.

[26] R. Zhao, W. Ouyang, H. Li, and X. Wang, "Saliency detection by multi-context deep learning," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2015, pp. 1265–1274.

[27] L. Wang, H. Lu, X. Ruan, and M.-H. Yang, "Deep networks for saliency detection via local estimation and global search," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, June 2015.

[28] G. Li and Y. Yu, "Visual saliency based on multiscale deep features," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, June 2015.

[29] Y. Li, X. Hou, C. Koch, J. M. Rehg, and A. L. Yuille, "The secrets of salient object segmentation," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2014.

[30] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. C. Berg, and L. Fei-Fei, "ImageNet Large Scale Visual Recognition Challenge," *Int. J. Comput. Vision*, 2015.

[31] A. Bergamo, L. Bazzani, D. Anguelov, and L. Torresani, "Self-taught object localization with deep networks," *arXiv preprint arXiv:1409.3964*, 2014.

[32] M. Oquab, L. Bottou, I. Laptev, and J. Sivic, "Is object localization for free?-weakly-supervised learning with convolutional neural networks," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2015, pp. 685–694.

[33] R. G. Cinbis, J. Verbeek, and C. Schmid, "Weakly supervised object localization with multi-fold multiple instance learning," *IEEE Trans. Pattern Anal. Mach. Intell.*, 2016.

[34] M. Oquab, L. Bottou, I. Laptev, and J. Sivic, "Learning and transferring mid-level image representations using convolutional neural networks," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2014, pp. 1717–1724.

[35] Z. Bolei, K. Aditya, L. Agata, O. Aude, and T. Antonio, "Learning deep features for discriminative localization," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2016.

[36] H. Bilen and A. Vedaldi, "Weakly supervised deep detection networks," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2016.

[37] J. R. Uijlings, K. E. van de Sande, T. Gevers, and A. W. Smeulders, "Selective search for object recognition," *Int. J. Comput. Vision*, vol. 104, no. 2, pp. 154–171, 2013.

[38] C. L. Zitnick and P. Dollár, "Edge boxes: Locating object proposals from edges," in *Proc. Eur. Conf. Comput. Vis.*, 2014, pp. 391–405.

[39] C. Sun, M. Paluri, R. Collobert, R. Nevatia, and L. Bourdev, "Pronet: Learning to propose object-specific boxes for cascaded neural networks," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2016.

[40] D. Pathak, P. Krahenbuhl, and T. Darrell, "Constrained convolutional neural networks for weakly supervised segmentation," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2015, pp. 1796–1804.

[41] R. Quan, J. Han, D. Zhang, and F. Nie, "Object co-segmentation via graph optimized-flexible manifold ranking," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, June 2016.

[42] D. Zhang, J. Han, C. Li, J. Wang, and X. Li, "Detection of co-salient objects by looking deep and wide," *Int. J. Comput. Vision*, pp. 1–18, 2016.

[43] M. D. Zeiler and R. Fergus, "Visualizing and understanding convolutional networks," in *Proc. Eur. Conf. Comput. Vis.*, 2014, pp. 818–833.

[44] A. Mahendran and A. Vedaldi, "Understanding deep image representations by inverting them," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2015, pp. 5188–5196.

[45] A. Dosovitskiy and T. Brox, "Inverting convolutional networks with convolutional networks," *CoRR abs/1506.02753*, 2015.

[46] B. Zhou, A. Khosla, A. Lapedriza, A. Oliva, and A. Torralba, "Object detectors emerge in deep scene cnns," *arXiv preprint arXiv:1412.6856*, 2014.

[47] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," *CoRR*, vol. abs/1409.1556, 2014.

[48] K. Chatfield, K. Simonyan, A. Vedaldi, and A. Zisserman, "Return of the devil in the details: Delving deep into convolutional nets," in *Proc. British Mach. Vis. Conf.*, 2014.

[49] K. He, X. Zhang, S. Ren, and J. Sun, "Spatial pyramid pooling in deep convolutional networks for visual recognition," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 37, no. 9, pp. 1904–1916, 2015.

[50] S. Lazebnik, C. Schmid, and J. Ponce, "Beyond bags of features: Spatial pyramid matching for recognizing natural scene categories," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2006.

[51] A. Vedaldi and B. Fulkerson, "VLFeat: An open and portable library of computer vision algorithms," http://www.vlfeat.org/, 2008.

[52] S. Shalev-Shwartz and T. Zhang, "Stochastic dual coordinate ascent methods for regularized loss minimization," *J. Mach. Learn. Res.*, vol. 14, no. Feb, pp. 567–599, 2013.

[53] R. Achanta, A. Shaji, K. Smith, A. Lucchi, P. Fua, and S. Susstrunk, "Slic superpixels compared to state-of-the-art superpixel methods," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 34, no. 11, pp. 2274–2281, 2012.

[54] M. Everingham, L. Van Gool, C. K. I. Williams, J. Winn, and A. Zisserman, "The PASCAL Visual Object Classes Challenge 2012 (VOC2012) Results," http://www.pascal-network.org/challenges/VOC/voc2012/workshop/index.html.

[55] E. Ahmed, S. Cohen, and B. Price, "Semantic object selection," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2014.

[56] C. Rother, V. Kolmogorov, and A. Blake, "Grabcut: Interactive foreground extraction using iterated graph cuts," *ACM Trans. Graphics*, vol. 23, no. 3, pp. 309–314, 2004.

[57] M. Everingham, L. Van Gool, C. K. I. Williams, J. Winn, and A. Zisserman, "The PASCAL Visual Object Classes Challenge 2007 (VOC2007) Results," http://www.pascal-network.org/challenges/VOC/voc2007/workshop/index.html.

[58] A. Opelt, A. Pinz, M. Fussenegger, and P. Auer, "Generic object recognition with boosting," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 28, no. 3, pp. 416–431, 2006.

[59] M. Rubinstein, A. Joulin, J. Kopf, and C. Liu, "Unsupervised joint object discovery and segmentation in internet images," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2013.

[60] M. Everingham, L. Van Gool, C. K. I. Williams, J. Winn, and A. Zisserman, "The PASCAL Visual Object Classes Challenge 2010 (VOC2010) Results," http://www.pascal-network.org/challenges/VOC/voc2010/workshop/index.html.

[61] P. Jiang, N. Vasconcelos, and J. Peng, "Generic promotion of diffusion-based salient object detection," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2015, pp. 217–225.

[62] Ç. Aytekin, E. C. Ozan, S. Kiranyaz, and M. Gabbouj, "Visual saliency by extended quantum cuts," in *Proc. Int. Conf. Image Proc.*, 2015, pp. 1692–1696.

[63] M. Marcin and S. Cordelia, "Accurate object recognition with shape masks," *Int. J. Comput. Vision*, vol. 97, no. 2, pp. 191–209, 2012.

[64] F. Brian, V. Andrea, and S. Stefano, "Class segmentation and object localization with superpixel neighborhoods," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2009.

[65] D. Aldavert, A. Ramisa, R. L. de Mantaras, and R. Toledo, "Fast and robust object segmentation with the integral linear classifier," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2010.

[66] Khan, Nazar, Tappen, and M. F, "Discriminative dictionary learning with spatial priors." in *Proc. Int. Conf. Image Proc.*, 2013.

[67] D. Pathak, E. Shelhamer, J. Long, and T. Darrell, "Fully convolutional multi-class multiple instance learning," in *Proc. ICLR*, 2015.

[68] G. Papandreou, L.-C. Chen, K. Murphy, and A. L. Yuille, "Weakly-and semi-supervised learning of a dcnn for semantic image segmentation," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2015.

[69] B. Ali, C. Ming-Ming, J. Huaizu, and L. Jia, "Salient object detection: A benchmark," *IEEE Trans. Image Process.*, vol. 24, no. 12, pp. 5706–5722, 2015.

[70] C. Yang, L. Zhang, H. Lu, X. Ruan, and M.-H. Yang, "Saliency detection via graph-based manifold ranking," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2013, pp. 3166–3173.

[71] B. Hariharan, P. Arbeláez, L. Bourdev, S. Maji, and J. Malik, "Semantic contours from inverse detectors," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2011, pp. 991–998.

[72] A. Joulin, F. Bach, and J. Ponce, "Discriminative clustering for image co-segmentation," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2010.

[73] J. Armand, B. Francis, and P. Jean, "Multi-class cosegmentation," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2012.

[74] G. Kim, E. P. Xing, L. Fei-Fei, and T. Kanade, "Distributed cosegmentation via submodular optimization on anisotropic diffusion," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2011.

[75] K. R. Jerripothula, J. Cai, and J. Yuan, "Object detection with discriminatively trained part-based models," *IEEE Trans. Multimedia*, 2016.

[76] R. Girshick, "Fast r-cnn," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2015, pp. 1440–1448.

[77] S. Ren, K. He, R. Girshick, and J. Sun, "Faster r-cnn: Towards real-time object detection with region proposal networks," in *Proc. Adv. Neural Inf. Process. Syst.*, 2015, pp. 91–99.

[78] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, "You only look once: Unified, real-time object detection," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2016.

Hisham Cholakkal received the Bachelor of Technology degree in electronics and communication engineering from Calicut University, India, in 2006 and the Master of Technology degree in signal processing from Indian Institute of Technology Guwahati, India in 2009. From 2009 to 2012 he worked at Central Research Lab of Bharat Electronic Limited, Bangalore, India and Advanced Digital Sciences Center, Singapore. He is currently pursuing the PhD degree with the School of Computer Science and Engineering, Nanyang Technological University, Singapore. His current research interests includes computer vision, image processing and machine learning.

Jubin Johnson received the Bachelor of Technology degree in electronics and communication engineering from Vellore Institute of Technology, India in 2010. From 2010 to 2012, he worked at Wipro Technologies, India. He is currently pursuing the PhD degree with the School of Computer Science and Engineering, Nanyang Technological University, Singapore. His current research interests include computer vision, image and video processing, and computer graphics

Deepu Rajan received the Bachelor of Engineering degree in electronics and communication engineering from the Birla Institute of Technology, Ranchi, India, the M.S. degree in electrical engineering from Clemson University, Clemson, SC, USA, and the Ph.D. degree from the Indian Institute of Technology Bombay, Mumbai, India. He is an Associate Professor with the School of Computer Science and Engineering, Nanyang Technological University, Singapore. From 1992 to 2002, he was a Lecturer with the Department of Electronics, Cochin University of Science and Technology, Cochin, India. His current research interests include image processing, computer vision, and multimedia signal processing.